LABOR PRODUCTIVITY: STRUCTURAL CHANGE AND CYCLICAL DYNAMICS

By

Martin Neil Baily*
McKinsey Global Institute

Eric J. Bartelsman**
Free University of Amsterdam

John Haltiwanger***
University of Maryland
NBER
Bureau of the Census

CES 98-7 May 1998

Abstract

A longstanding issue in empirical economics is the behavior of average labor productivity over the business cycle. This paper provides new insights into the cyclicality of aggregate labor productivity by examining the cyclical behavior of productivity at the plant level as well as the role of reallocation across plants over the cycle. We find that plant-level productivity is even more procyclical than aggregate productivity because short-run reallocation yields a countercyclical contribution to labor productivity. At the plant level we find that cyclicality of productivity varies systematically with long-run employment growth. Over the course of the cycle, plants that are long-run downsizers exhibit significantly greater procyclicality of productivity than long-run upsizers. When we control for the direction of a cyclical shock, we find that the fall in productivity from an adverse cyclical shock for long-run downsizers is significantly larger in magnitude than the fall in productivity from an equivalent adverse cyclical shock for long-run upsizers. We argue that these findings raise questions about one of the most popular explanations or procyclical productivity: changing factor utilization over the cycle.

Keywords: Cyclical Productivity

*McKinsey Global Institute; **Free University of Amsterdam; *** University of Maryland, NBER and Census Bureau. The views expressed herein are solely the authors'. We wish to thank the staff and visiting researchers at the Center for Economic Studies, U.S. Bureau of the Census, for providing a stimulating research environment. We would also like to thank the participants of the NBER Summer '93 workshop, the 1995 NBER Small Group Meeting on Alternative Explanations of Employment Fluctuations, the 1995 STE workshop, the Econometric Society meetings, January 1995, and seminar participants at several universities and institutions for valuable comments and suggestions. Lucia Foster provided superb research assistance.

1 Introduction

A longstanding issue in empirical economics is the behavior of average labor productivity over the business cycle. While measurement difficulties loom large, the preponderance of evidence suggests that average labor productivity falls in recessions and increases in booms.[1] There have been no shortage of explanations of the behavior of labor productivity over the business cycle.[2] Accounting for the cyclical behavior of labor productivity has, traditionally, fallen into three categories: the underlying theoretical framework is adjusted, by assuming market imperfections or non-standard production technology, the data are adjusted for possible measurement error due, say, to unmeasured changes in factor utlization, or the primary driving force for cycles is assumed to be technology shocks. In our view, there is an inherent limitation in most of the existing empirical analysis given that the typical empirical study is based on aggregate (e.g., total economy or industry-level) data.[3] In this paper, we make a step towards remedying this limitation by examining the cyclical behavior of productivity for a large sample of plants for the 1970s and the 1980s.

Our plant-level analysis permits us to distinguish the cyclical behavior of aggregate labor productivity from plant-level productivity. The cyclical behavior for the average plant can differ from the aggregate owing to changing shares across plants over the business cycle and to differing levels of productivity across plants. In our empirical analysis, we decompose the change in aggregate productivity into the contribution of within plant and between plant components for continuing plants as well as the net contribution to productivity of plant entry and exit. We find that the within plant component for continuing plants exhibits greater procylicality than aggregate labor productivity. This finding reflects the countercyclical pattern exhibited by the between plant component. That is, we find that labor shares of less productive plants falls in recessions, thereby

[1] See, e.g., Cooley and Prescott [6] who report that the correlation between quarterly detrended output per hour and output is 0.34 indicating mildly procyclical labor productivity. As one referee put it, this mild procyclicality is especially puzzling in light of the strongly countercyclical variation in the measured capital labor ratio which would suggest strongly countercyclical labor productivity.

[2] This is a long literature with antecedents in the Dunlop/Tarshis/Keynes debate about the behavior of real wages over the business cycle. Notable contributions to this literature across the decades include Kuh [16], Fair [9], Solow [18], Fay and Medoff [10], Hall [14] and [15], Gordon [12], Bernanke and Parkinson [4], Burnside, Eichenbaum, and Rebelo [5].

[3] A notable exception is Aizcorbe [1], who analyzes cyclical labor productivity in automobile assembly plants. In auto assemblies retaining workers not needed for current production turns out to prevalent, likely owing to the specific nature of labor contracts in that industry. Industry-experts have joked that in the 1980s unionized auto workers, retained under a "no-layoff" contract, painted every church in Michigan.

reducing the procyclicality of aggregate labor productivity.

Having established that the procyclicality of labor productivity is a within plant phenomenon, we then proceed to investigate the cyclical patterns of labor productivity at the plant level. A key aspect of our empirical approach is to explore the relationship between heterogeneity in *long-run* structural changes across individual plants and heterogeneity in patterns of *short-run* procyclical productivity. The motivation for this approach is based upon an indirect test of one of the primary explanations of changes in labor productivity over the cycle: changing factor utilization rates over the cycle. As is commonly argued, if factor utilization rates are lower in recessions but are not properly measured, then output will fall more than measured factor inputs (in particular, more than the fall in labor input), resulting in procyclical productivity. We propose an indirect test of this hypothesis by observing that the incentives for an individual establishment to vary the utilization rate of a factor as opposed to the level of factor input itself depends on the future conditions anticipated by the establishment.

A simple example illustrates the argument and, in turn, the nature of our empirical approach. Consider the differing incentives of an establishment that is anticipating a decrease in employment over the long run relative to an establishment anticipating an increase in employment over the long run. During a downturn from an adverse demand shock, the long-run upsizing plant will be more reluctant to lay off workers since the plant is anticipating higher employment needs in the future. Accordingly, in response to an adverse shock, the long-run upsizing plant is more likely to lower the utilization rate of labor (or other inputs) in order to retain workers and thus a long-run upsizing plant is more likely to exhibit a decrease in measured labor productivity. Note, however, that the argument is not symmetric – the opposite implications hold for a positive demand shock. During a boom, a long-run downsizing plant will be more reluctant to add workers since the plant is anticipating lower employment needs in the future. Accordingly, in response to a positive demand shock, a long-run downsizing plant will be more likely to increase the utilization rate and thus should exhibit a greater increase in measured productivity. These arguments suggest that we should observe systematic and different patterns in the cyclicality of productivity between long-run upsizing and downsizing plants and, in turn, these differences should depend on the direction of the cyclical shock.

In our empirical analysis, we track individual plants over an extended period of time and classify them as long-run upsizers or downsizers accordingly (on a peak-to-peak basis over the cycle). This classification by long-run changes is relatively straightforward and we indeed find systematic differences in the behavior of cyclical productivity across plants classified in this manner. By itself, such systematic differences by ex-post long-run changes does not constitute a test of the above discussed hypothesis. Two additional issues need to be addressed. First, it is important to isolate the fluctuations in plant-level productivity due to demand shocks and to control for the magnitude of the shocks. We do this by exploiting the downstream demand indicators developed by Bartelsman, Caballero and Lyons [3]. We extend their approach by using Shea [17] type conditions to restrict the downstream indicator to reflect those components that are arguably exogenous. Secondly, it may be inappropriate to use ex-post indicators of long-run changes since the patterns of short-run cyclical productivity may influence the indicator of long-run change. Moreover, the argument above suggests that it is the anticipated future conditions that matter. To correct for both the endogeneity problem and the related need to use anticipated changes, in some parts of our empirical analysis we use ex-ante forecasts of long-run changes to classify plants as long-run upsizers and downsizers.

The organization of our paper is as follows: first we give a description of the data and basic facts. Second, we provide a formal empirical decomposition of aggregate changes in labor productivity into within plant, between plant and net entry components. Third, we explore the cyclical patterns of labor productivity at the plant level and their relationship to long-run employment changes. The last section concludes with a brief summary.

2 Data Description and Exploration

2.1 Data Description

In this study we make use of the Longitudinal Research Database (LRD), available at the Bureau of the Census. The analysis uses the Annual Survey of Manufactures (ASM) portion of the LRD for the years 1972 through 1989. For our aggregate decompositions, we consider both the entire ASM which is a representative sample of manufacturing plants (see, Davis, Haltiwanger and Schuh [8] for more details) and a balanced panel of plants consisting of those plants which were in continuous operation from 1972 through 1989 (and had positive employment and shipments in all years).[4] The use of the entire ASM permits us to assess the contribution of entry and exit to the cyclical behavior of productivity as well as a comprehensive analysis of the impact of the reallocation of labor across plants across years. It turns out that the balanced panel, for the most part, yields similar results which provides confidence for our use of the balanced panel in our subsequent plant-level econometric analysis.

We measure output, labor, and productivity as follows. For our aggregate decompositions, output is measured as deflated shipments using 4-digit industry benchmark-years weighted output deflators from the NBER productivity dataset.[5] Labor input is measured as total hours for the year. In the ASM, total hours for production workers and total number of nonproduction workers are measured. Following methodology developed by Davis and Haltiwanger [7], we multiply the reported number of nonproduction workers with an estimate of the average hours per nonproduction worker to generate a measure of total nonproduction worker hours at the plant level. The estimate of average hours per nonproduction worker is based upon tabulations by 2-digit industry from the Current Population Survey (CPS) for each year. The measure of total hours at the plant is then simply the sum of production and nonproduction total hours. For some of the empirical analysis, we also measure labor as simply total employment with an accompanying output per worker measure of labor productivity.

Figure 1 provides a comparison of the output, labor and productivity growth rates from the entire ASM (all LRD using appropriate ASM sample weights) and from our balanced panel. The top panel shows output and labor input for the entire ASM, the middle panel displays the same series for our balanced panel, while the bottom panel compares output per hour from the entire ASM and our balanced sample. The central point of Figure 1 is that the basic aggregate properties of output, employment, and productivity fluctuations in the balanced panel are essentially the same

[4] Although the plant-level data for the ASM are now available through 1993, our analysis stops in 1989 for the following reasons: (1) a 1989-93 subperiod would be quite short and not reflect a peak-to-peak subperiod as for the other subperiods; (2) the longitudinal linkages at the plant-level for the 1990s are still being improved and we would accordingly have less confidence in any results based upon longitudinal data for the 1990s.

[5] We deliberately chose to use the simplest measure of gross output when working with the entire LRD. It is possible to adjust shipments for the change in inventories but our confidence in this adjustment is much greater for the balanced panel than for the entire ASM since imputation rates for inventories are high for small plants. Moreover, it is possible to consider a value added per worker measure but again our confidence in such a measure is much greater for the balanced panel for similar reasons. In the subsequent econometric analysis using the balanced panel, we consider both of these alternative measures of output.

as in the all-ASM panel.[6]

Several basic patterns are evident from Figure 1: total hours are procyclical. output is more cyclically volatile than total hours and, accordingly, average labor productivity is procyclical. In the all LRD panel, the standard deviation of the growth rate of output is 5.6, the standard deviation of the growth rate of labor is 5.2, and the correlation between labor productivity and output is 0.29. For the balanced panel, the same statistics are 5.6, 4.6 and 0.63. Thus, the balanced panel exhibits greater procyclicality of labor than the entire ASM. As we will see in the next section, part of this difference reflects composition or reallocation effects rather than differences in the behavior of the average plant.

2.2 The Decomposition of Aggregate Productitivty Growth into Within, Between and Net Entry Effects

We now turn to an exploration of reallocation or composition effects to determine the extent which the cyclical patterns of aggregate productivity occur within plants or are the result of shifts in the allocation of labor across plants, including entry and exit.

We can decompose annual productivity changes into a within plant component, a between plant component, and a net entry term.[7] This last term is necessarily zero for the balanced panel of continuing plants.

$$(1)\ \Delta\Pi_t = \sum_{i\in Cont} \overline{\phi_i}\Delta\Pi_{it} + \sum_{i\in Cont} \Delta\phi_{it}(\overline{\Pi_i} - \overline{\Pi}) + \sum_{i\in Entry} \phi_{it}(\Pi_{it} - \overline{\Pi}) - \sum_{i\in Exit} \phi_{t-1,i}(\Pi_{i,t-1} - \overline{\Pi}),$$

[6] In unreported results, we have also examined the relationship between the all ASM sample and the series from the published ASM (e.g., the NBER productivity database). Note that there are some modest differences between the published ASM and the weighted sample ASM results. The differences reflect some technical adjustments (the fixed based difference) and the influence of very small (less than 5 employee establishments) that are excluded from the mail universe of the ASM (see the technical appendix of Davis, Haltiwanger and Schuh [8] for further discussion).

[7] The decomposition we employ is a modified version of that developed by Griliches and Regev [13]. In prior work (and a previous version of this paper), we have considered a more complex decomposition for continuers that involves a within plant, a between plant term and a cross term. The alternative decomposition defines the within term using base period employment shares, the between plant term uses base period productivity, and the cross term involves the sum of the cross products of share and productivity changes. This more complex decomposition may be conceptually preferred. For example, in this alternative decomposition, the within plant term can be readily interpreted as the change in productivity that would have occurred if shares had stayed at their initial values. However, this alternative decomposition is more subject to measurement error (see, Foster, Haltiwanger and Krizan [11] for further discussion). Moreover, in the current context, the simpler Griliches-Regev type decomposition yields very similar results to those generated by the more complex alternative decomposition.

where Π_{it} is output per hour in plant i at time t, Π_t is aggregate output per hour at time t, ϕ_{it} is the share of hours at plant i at time t, and a bar over a variable indicates the average of the variable across $t-1$ and t. The first term in this decomposition is a within plant component reflecting the change in productivity from continuing plants holding labor shares at their time series average. The second term in this decomposition is a between plant component reflecting the change in labor share from continuing plants for fixed levels of productivity. The plant-level productivity in this term is deviated from the overall average so that an increase in the share of labor for a plant contributes positively only if the plant has higher than the overall average productivity.[8]

The last two terms reflect the contribution of net entry. Births contribute positively to the change to the extent that they enter above the average and deaths contribute positively to the change to the extent that they exit lower than the average. In what follows, we convert this decomposition to rates by dividing each of the terms by the aggregate level of productivity in $t-1$.

The upper panel shows the results of the decomposition using the full LRD. Since the decomposition requires measuring plant-level changes, the results reported in the upper panel exclude 1974, 1979, and 1984 which are the first years of new ASM panels. As discussed in detail in Davis, Haltiwanger and Schuh [8], it is only possible to link large, continuing plants across ASM panels. In the figure, the values for 1974, 1979, and 1984 are interpolated so appropriate caution should be used in those years. For the balanced panel of continuing plants (lower panel) we can examine the cyclical behavior of the components of the decomposition.

In the upper panel, the within plant term exhibits pronounced procyclical behavior, with sharp increases in the recovery years 1976 and 1983. The correlation of the within plant component with aggregate output growth (excluding first ASM panel years) is 0.63. Recall from Figure 1 that the correlation between aggregate productivity and output growth is 0.29. Thus, the within plant component exhibits a greater degree of procyclicality than the aggregate measure. The reason for this is evident from the behavior of the between plant component. The correlation of the between plant component with aggregate output growth is -0.31. As is evident from the figure, during recessions the between plant component increases. This finding reflects the shift of labor shares away from less productive and towards more productive plants. Thus, there is a countercyclical

[8] In a balanced panel, it is unnecessary to deviate from the overall mean explicitly since the sum of the change in shares is equal to zero. This property does not hold for continuing plants in an unbalanced panel.

tendency in the behavior of aggregate productivity reflecting the lower employment shares of less productive plants during recessions.

The contribution of net entry to the annual change in aggregate productivity is modest. This result, which contrasts with recent results in the literature (see, e.g., Baily, Bartelsman, and Haltiwanger [2] and Foster, Haltiwanger and Krizan [11]) is actually not surprising. In that literature, the important positive contribution of net entry to aggregate productivity growth occurs over a longer (five- or ten-year) time horizon. In an accounting sense, this is because the share of labor accounted for by entering and exiting plants for annual changes is small while the share of labor accounted for by entering and exiting plants for longer-run changes is large. For example, for annual changes, the share of labor accounted for by exiting plants is 0.02 and for entering plants is 0.01. However, for 10-year changes, the share accounted for by entrants in year t is 0.26 while the share accounted for by exits in year t-10 is 0.28. Moreover, the productivity differential between entering and exiting plants increases with the horizon over which the changes are measured due to learning and selection effects. Thus, while the contribution of net entry to productivity growth is modest at an annual frequency, it becomes quite large even at five-year intervals. Foster, Haltiwanger and Krizan [11] report that the contribution of net entry for five-year changes to aggregate productivity exceeds 20 percent.

It seems that a fair general characterization of the results is that procyclical productivity is a within-plant phenomenon. Put differently, it is the way that plant managers adjust operations within their plants in response to shocks that determines the overall cyclical response of the manufacturing sector. Since it is a within-plant phenomenon, this enables us in subsequent sections to pursue a simple regression analysis of within-plant changes in productivity over the cycle. In this regard, it is important to emphasize that even though the cyclical variation in productivity is primarily due to within-plant effects, this does not imply that there are no important cross sectional differences in the procyclicality of productivity. In particular, we are interested in whether there are differences in the within-plant changes in productivity over the cycle between plants with different long-run outlooks. We now turn to our investigation of such differences.

2.3 Basic Facts: Long-run Structural Change and Cyclical Dynamics

To begin the exploration of the connection between the cyclical variations in productivity and the long-run structural changes, we compare the cyclical productivity fluctuations of plants that experienced increases in employment over the long run to those that decreased their employment. By long run in this context, we refer to changes that correspond roughly to peak-to-peak changes over our sample. In particular, we focus on peak-to-peak changes from 1973-79 and from 1979-89. Figure 3 shows the time series pattern of average employment and labor productivity growth for long-run upsizers and long-run downsizers for our sample of plants. Table 1 presents summary statistics for these growth rates for plants classified on these dimensions. For this descriptive exercise, the growth rates are measured as log changes at the plant level and we calculate the average across plants in each year (cross classified by an indicator of long-run change). Long-run upsizers (downsizers) for, say, 1973-79 are those plants that had higher (lower) employment in 1979 than 1973. Plants that downsized employment over the long run exhibit greater volatility in the growth rate of productivity than do plants that are long-run upsizers.

This pattern is suggestive that differences in long-run structural changes may be important in understanding cyclical variations in productivity. However, it is only an illustrative exercise which in part highlights several factors that we need to take into account in our more formal empirical analysis in the next section before any conclusions can be drawn. First, the observed pattern might simply result from the fact that the two different plant groups are hit by shocks of different magnitudes. The different cyclical patterns of employment growth may reflect such differences in the shock patterns and magnitudes. This could occur, for example, because of differences in the specific sectoral composition within each group.

Second, as we emphasized in the introduction, the cyclical behavior of productivity for long-run upsizers and downsizers is apt to vary by the sign of the shocks. Recall the argument is that long-run upsizers should be reluctant to layoff workers in response to negative shocks but that long-run downsizers should be reluctant to hire workers in response to positive shocks.

Third, the allocation of plants into upsizers and downsizers is endogenous. For example, plants that suffer a severe downward short-term shock may be more likely to end up as long-term downsizers. And, similarly, plants that experience a large upward short-term shock may be more likely to be long-run upsizers. To deal with this problem we check the robustness of our conclusions when plants are allocated into the two groups on the basis of ex-ante forecasts of their long-run status.

Finally there is an additional empirical issue that we tackle. Dividing plants into upsizers and downsizers allows a useful way of thinking about the long-run behavior that will affect short-run procyclicality. But it is of course arbitrary to draw a hard line at the zero level of long-run employment change. The short-run employment changes can be expected to change continuously with expected long-run employment needs. We therefore test an econometric specification where short-run behavior interacts with the magnitude of employment change in the long run.

3 Empirical Specification

The focus of our empirical analysis is to estimate simple regressions relating plant-level productivity to indicators of cyclical shocks. We begin with the following specification:

$$\hat{\pi}_{it} = \pi_{it} - \hat{a}_i - \hat{b}_i t = \sum_n \gamma_n I_{ni} D_{it} + \varepsilon_{it} \tag{2}$$

where π_{it} is the log of plant productivity, I_{ni} is an indicator variable indicating whether the plant is an upsizer or downsizer, and D_{it} is an index of downstream demand. Thus, our dependent variable is log labor productivity measured as a deviation from a plant-specific linear trend.[9] The cyclical variation in plant-level productivity deviations is allowed to depend on whether the plant is a long-run upsizer or a long-run downsizer–captured by the variable I_{ni} on the right hand side of equation (2). We measure the size of the shock facing the plant, denoted by the variable D_{it} (in some later specifications separating out the positive and negative values). The parameter γ_n is the estimated parameter whose size reflects the extent to which plant level productivity varies with short-run shocks. Equation (2), therefore, permits a test of the hypotheses that 'cyclical' variations in plant level productivity depend upon whether the plant is increasing or decreasing its level of employment over the long run. We also consider two modifications of this basic specification.

In some specifications, we permit the response to depend on the magnitude of the long-run

[9]The basic measure of output used is gross output (shipments adjusted for inventory change deflated by the four-digit industry deflator for the industry in which the plant is located). Results using value added are not reported in this paper but they show no substantial difference from the reported results. The cyclical response of labor productivity when calculated using employees is somewhat different from the response when calculated with hours, as would be expected, and we report our results both ways.

employment change. Specifically, we consider the following specification:

$$\hat{\pi}_{it} = \sum_n \gamma_n I_{ni} D_{it} |\Delta L_i| + \varepsilon_{it} \tag{3}$$

where $|\Delta L_i|$ is the absolute value of the change in (log) employment over the long run.

In some specifications, we interact the demand indicator with whether the demand change is positive or negative to permit the the response to a positive demand shock to be different than a negative demand shock. Relative to equation (2), this yields:

$$\hat{\pi}_{it} = \sum_n \sum_d \gamma_{nd} I_{ni} D_{it} I_{dit} + \varepsilon_{it} \tag{4}$$

where I_{dit} is a dummy variable indicating whether the change in the demand indicator is positive or negative and γ_{nd} varies by both the long-run change and the direction of the shock. Finally, we consider the most general interacted specification given by:

$$\hat{\pi}_{it} = \sum_n \sum_d \gamma_{nd} I_{ni} D_{it} I_{dit} |\Delta L_i| + \varepsilon_{it}. \tag{5}$$

Before proceeding to the results, the next subsections provide further clarification on how the right hand side variables in these alternative specifications are defined.

3.1 Classifying Plants as Long-run Upsizers or Downsizers

The simplest way to determine whether or not a plant is a long-run upsizer or downsizer is to see what actually happened to it over the long run (we consider 1973-79 and 1979-89 as two separate "long-run" periods, each reflecting a full business cycle). The results from this procedure are reported in the results in the next section as the "ex-post" estimates. An alternative approach is to base the assignment of the plants on an estimate of the (ex-ante) expectation of each plant's long-run employment status. Thus, the indicator I_{ni} will now depend on plant level expectations of long-run employment. We consider this latter approach only for the 1979-89 subperiod since we don't have the appropriate variables to forecast the changes for the 1973-79 subperiod.

In practice, the ex-ante projections are computed by regressing the ex-post realizations of these growth coefficients for the pooled observations in each 2-digit industry on a set of variables which are observable by the plant before 1979. These variables are: plant size (average employment in 1977 and 1978), plant wage (average compensation per employee in 1977 and 1978), lagged plant productivity (average labor productivity in 1976 and 1977), plant age, plant productivity and employment growth rates from 1972 through 1978, cumulative real investment from 1972 through 1978 relative to average plant size, and controls for 4-digit industry, region, and ownership type. The R^2 of these projection equations range from 0.1 to 0.65 across industries.

3.2 Measuring the Magnitude and Direction of the Shocks

The variable D_{it} is a downstream demand indicator specific to the 4-digit industry to which the plant is assigned. Our starting point here is the downstream activity indicator constructed by Bartelsman, Caballero and Lyons [3] (hereafter BCL). The BCL downstream indicator is a weighted average of changes in economic activity (measured by cost-share weighted aggregate of factor inputs) of other industries and service sectors, with weights equal to their share of purchases of the output from the industry in question. As argued by BCL, the downstream indicator isolates the component of procyclical productivity associated with demand fluctuations.

Shea [17] uses output in downstream sectors to serve as demand instruments. Shea provides two criteria which output from downstream sectors must satisfy in order to qualify as instruments, namely relevance and exogeneity. Relevance is satisfied when the downstream industry purchases intermediates that comprise a large portion of the upstream industry's output. Exogeneity is satisfied when the purchases from the upstream industry constitute only a small fraction of total material expenditures of the downstream industry.

Our extension of the BCL downstream indicator is to filter out downstream sectors which fail to meet the exogeneity criteria. Specifically, using input-output matrices and detailed gross output time series from the national accounts, we exclude from the BCL downstream indicator for each industry those downstream industries whose purchases of the upstream industry are larger than 5 percent of their expenditures on intermediate inputs. This rules out the endogenous effect that the upstream industries' productivity may have, via intermediate goods prices, on the downstream industry's activity. The BCL downstream indicators used in this paper meet the Shea-relevance criteria by construction, because they are weighted averages of activity in all (or all exogenous) downstream industries.

Because the Shea-exogeneity adjustment of the BCL indicator requires consistently matched input-output and national accounts data over time, we were able to readily undertake this adjustment only for the years 1977-89. Thus, we have the Shea exogeneity adjusted downstream indicators only for our analysis of the 1979-89 period. We use the previously constructed BCL downstream indicators for the 1973-79 period. Note that we obtained quite similar results for our 1979-89 analysis when we use the non-exogeneity adjusted BCL indicators as an alternative. Thus, at least in the current context, we found our downstream indicators to yield robust results regardless of whether or not we made the exogeneity adjustment.

The downstream indicators used in the regressions, D_{it}, are deviations from linear trend of the cumulative index of aggregated downstream activity changes as described above.

4 Empirical Findings

The results of our investigation are presented in tables 2-4.[10] Table 2 presents results of simply regressing plant level productivity (as measured by the LHS of equation (2) using either output per hour or output per worker) on the downstream indicator. The results indicate significant procyclicality of plant-level productivity for both the subperiods and for both measures of productivity.

In Table 3a, we present the results for 1979-89 using the ex-post classification of long-run upsizers and downsizers. The first two rows of the upper panel present the results when we permit the cyclicality to differ by the sign of the long-run change. The coefficient on long-run downsizers is positive and larger than that for long-run upsizers indicating that long-run downsizers exhibit greater procyclicality of productivity (measured either by output per hour or output per worker). Table 4 reports F-tests on whether the differences in coefficients are statistically significant and the results show that the differences are highly significant. The differences in coefficients are not only statistically significant but large in magnitude as well. For example, a shock of 10 percent yields a change in productivity of 3.46 percent for long-run downsizers and a change in productivity of only 0.66 percent for long-run upsizers.

The second two rows report the results when the response is permitted to vary continuously with the magnitude of the long-run change with a kink permitted depending on the direction of the

[10] We have also produced versions of the tables (not reported) with output measured as double-deflated value added instead of gross output. The results for value added are similar to those described in the text.

long-run change. The reported coefficients and the associated F-tests in Table 4 show that long-run downsizers exhibit significantly greater procyclicality of labor.

The lower panel repeats this same exercise but now separates out for each plant and for each year whether the plant faces a positive or negative demand shock (expansion or contraction). Positive procyclicality is again evident and the long-run downsizers show the greatest amount of procyclicality in both expansions and contractions. Moreover, the procyclicality for downsizers is greatest in response to negative shocks. The F-tests in Table 4 show that the differences between long-run upsizers and downsizers are highly statistically significant.

There is some evidence of procyclicality among the long-run upsizers when they are hit by adverse shocks, so we do not claim to rule out altogether the idea that plants retain workers in the short-run that they will need in the long run. Moreover, there is relatively little evidence of procyclicality among long-run upsizers in expansions which is consistent the view that long-run upsizers are quite willing to hire workers in response to a positive shock. Still, the long-run downsizers exhibit substantially greater procyclicality, especially during contractions. The results indicate that the elasticity of productivity with respect to a negative shock is always more than twice as large for long-run downsizers. Moreover, if we restrict our attention to output per hour results (presumably preferred in this context), the elasticity of productivity with respect to a negative shock is always more than three times larger for long-run downsizers.

Table 3b shows the analogous results using the forecast values of long-run employment changes to assign plants as upsizers and downsizers. The table shows that our findings are robust to the use of ex-ante classification of plants into long-run upsizers and downsizers. The largest procyclical effects are always with the long-run downsizers whether in expansions or contractions. In addition, the differences between long-run upsizers and downsizers are highly statistically significant (with one exception in which the difference is statistically different only at the 10 percent level). The magnitude of the difference is still very large when we use two distinct groups. The magnitude of the difference between long-run upsizers and downsizers is somewhat muted when we allow the impact to vary by the magnitude of the long-run changes and the direction of the shock (equation (5). However, this may reflect the fact that we are using ex-ante measurement and classification of long-run changes.

Table 3c shows the results for the 1973-79 subperiod. For these results, we are restricted to examining the ex-post classification of long run changes.[11] We again find that long-run downsizers exhibit greater procyclicality, both in expansions and in contractions. In addition, the differences between long-run upsizers and downsizers are highly statistically significant.

5 Conclusions

The contribution of this paper to the literature on cyclical productivity is to point to some clear empirical patterns in the data that are revealed at the plant level. At the aggregate level, procyclicality could occur if high productivity plants experienced larger cyclical variation and increased their share of employment in upturns and decreased it in downturns. Our findings show that reallocation works in the opposite direction. The impact of the reallocation of employment shares over the cycle yields a countercyclical component to aggregate productivity and thus the average plant, given fixed employment shares (the within plant component), exhibits greater procyclicality than aggregate productivity.

Not only do we find that procyclicality occurs within plants but we find that the cyclicality of plant level productivity disproportionately occurs among plants that are long-run downsizers. When such plants experience a negative short-run demand shock their productivity falls substantially. These empirical findings raise questions about one of the most popular explanations of changes in productivity over the cycle; namely, changes in factor utilization over the cycle that are not fully captured in the measured changes in inputs. If cyclical variation in utilization rates is driving cyclical fluctutations in productivity, then plants that are long-run upsizers should have a greater incentive to avoid layoffs and decrease factor utilization rates during downturns. Accordingly, long-run upsizers should exhibit a greater decrease in productivity during a contraction. However, we find just the opposite. That is, even controlling for the magnitude and direction of cyclical demand shocks, we find that long-run downsizers exhibit substantially larger elasticities of productivity with respect to negative demand shocks.

[11] In addition, we cannot make the Shea-type exogeneity adjustment for the long-run downsizers.

References

[1] Ana M. Aizcorbe. Procyclical Labour Productivity, Increasing Returns to Labour, and Labour Hoarding in Car Assembly Plant Employment. *Economic Journal*, 102(413):860–73, July 1992.

[2] Martin N. Baily, Eric J. Bartelsman, and John Haltiwanger. Downsizing and Productivity Growth: Myth or Reality? *Small Business Economics*, 8:259–278, 1992.

[3] Eric J. Bartelsman, Ricardo J. Caballero, and Richard K. Lyons. Customer and Supplier Driven Externalies. *American Economic Review*, 34(84-4):1075–84, September 1994.

[4] Ben Bernanke and Martin Parkinson. Procyclical Labor Productivity and Competing Theories of the Business Cycle: Some Evidence from Interwar U.S. Manufacturing Industries. *Journal of Political Economy*, 99(3):439–59, June 1991.

[5] Craig Burnside, Martin Eichenbaum, and Sergio Rebelo. Labor Hoarding and the Business Cycle. *Journal of Political Economy*, 101:245–73, April 1993.

[6] Thomas Cooley and Edward Prescott. Economic Growth and Business Cycles. In Thomas Cooley, editor, *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, 1994.

[7] Steven J. Davis and John Haltiwanger. Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986. In *Brookings Papers: Microeconomics*, Washington, D.C., 1991.

[8] Steven J. Davis, John Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. MIT Press, Cambridge, 1996.

[9] Ray C. Fair. *The Short-run Demand for Workers and Hours*. North-Holland, Amsterdam, 1969.

[10] J. Fay and James Medoff. Labor and Output Over the Business Cycle. *American Economic Review*, 75:638–55, 1985.

[11] Lucia Foster, John Haltiwanger, and C.J. Krizan. Aggregate Productivity Growth: Lessons from Microeconomic Evidence. Bureau of the Census, 1998.

Figure 1: Productivity and its Components

All LRD: Output and Labor Growth

Percent change

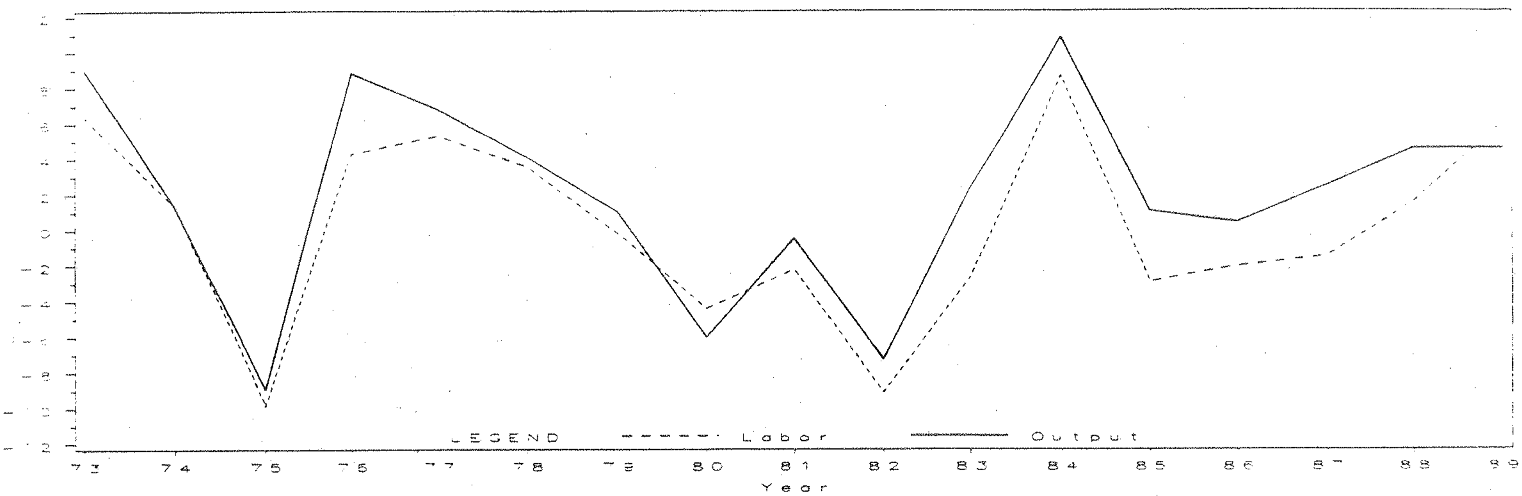

Balanced Panel: Output and Labor Growth

Percent change

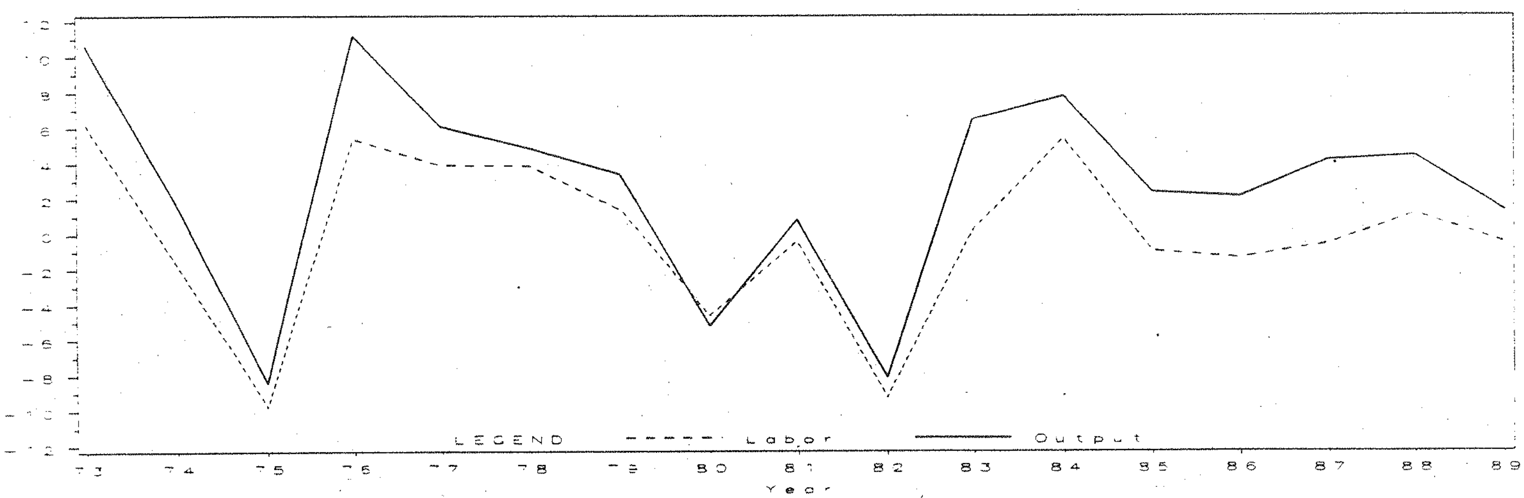

Productivity Growth Comparisons

Percent change

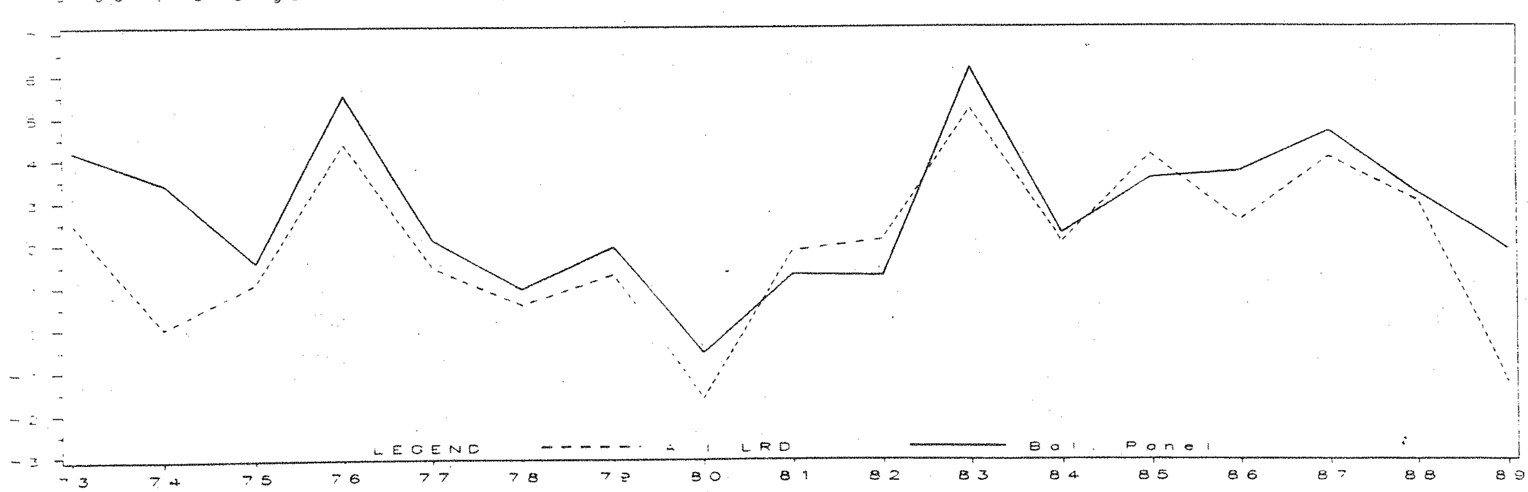

Source: Tabulations from the LRD

[12] Robert J. Gordon. Are Procyclical Productivity Fluctuations a Figment of Measurement Error. Northwestern University, 1990.

[13] Zvi Griliches and Haim Regev. Productivity and Firm Turnover in Israeli Industry: 1979-1988. *Journal of Econometrics*, 65:175–203, 1995.

[14] Robert E. Hall. The Relation between Price and Marginal Cost in U.S. Industry. *Journal of Political Economy*, 96:921–47, October 1988.

[15] Robert E. Hall. Econometric Research on Shifts of Production Functions at Medium and High Frequencies. presented at Barcelona World Congress of the Econometric Society, August 1990.

[16] Edwin Kuh. The Brookings Quarterly Econometric Model of the United States. In James Duesenberry, Gary Fromm, Lawrence Klein, and Edwin Kuh, editors, *Income Distribution and Employment Over the Business Cycle*. Rand McNally, Chicago, 1965.

[17] John Shea. Do Supply Curves Slope Up? *Quarterly Journal of Economics*, 108(1):1–32, 1993.

[18] Robert M. Solow. Some Evidence on the Short-Run Productivity Puzzle. In J. Baghwati and Eckaus R.S., editors, *Development and Planning*, pages 317–25. Allen and Unwin, London, 1973.

Figure 2: Productivity Decompositions

All LRD

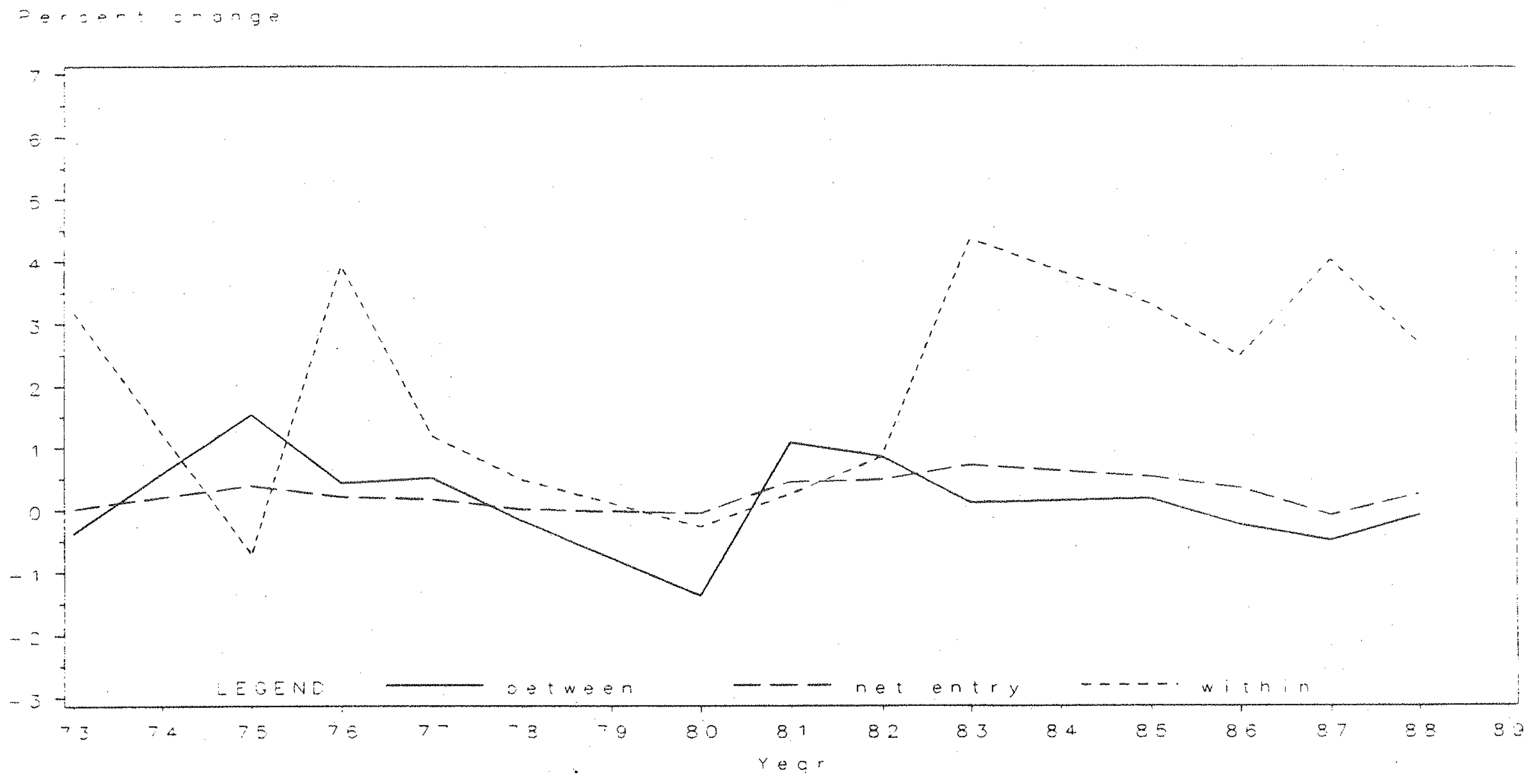

Balanced Panel

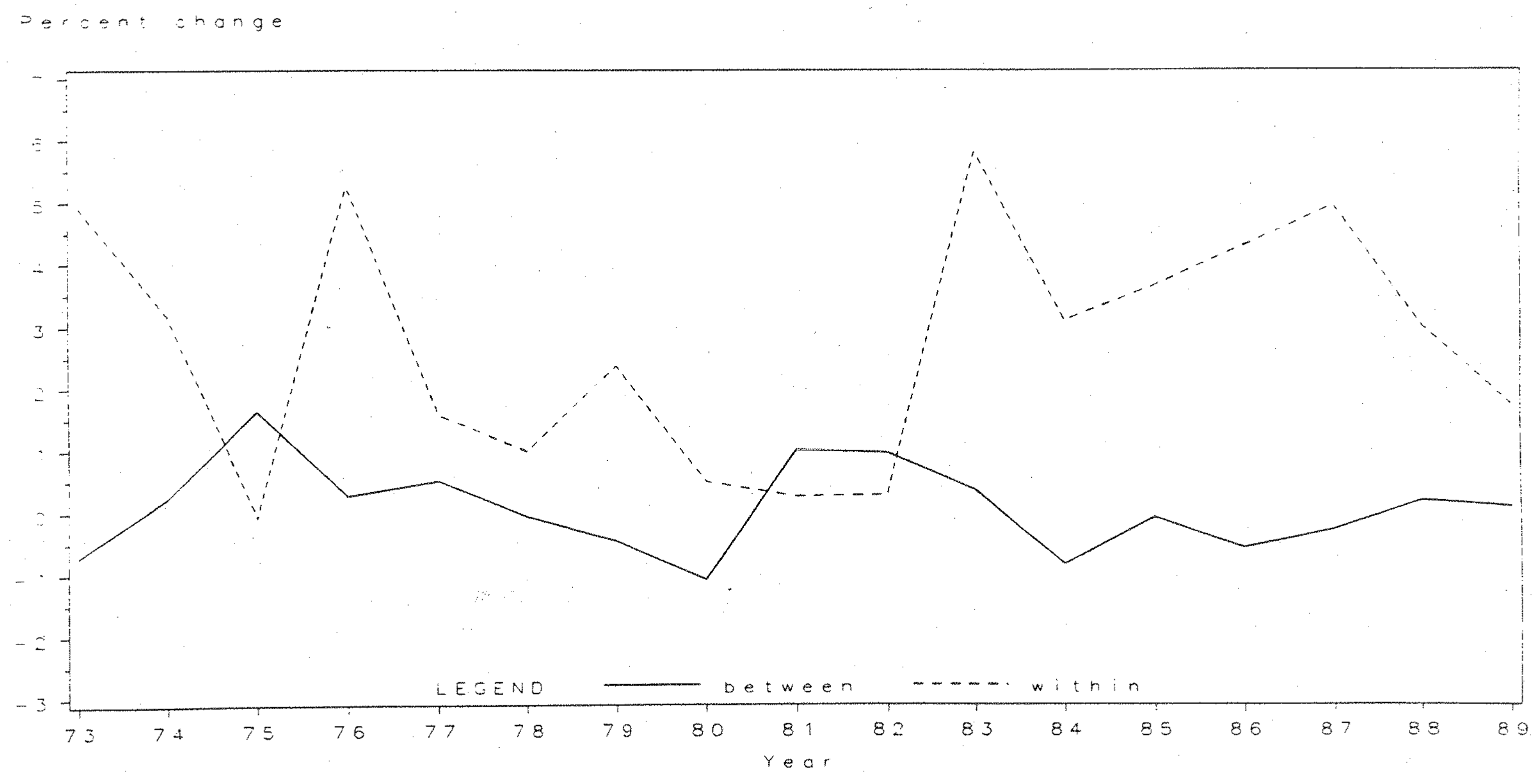

Source: Tabulations from the LRD.
Note. All LRD results are interpolated for 1974, 1979, and 1984.

Figure 3. Productivity and Labor Growth for Upsizers and Downsizers

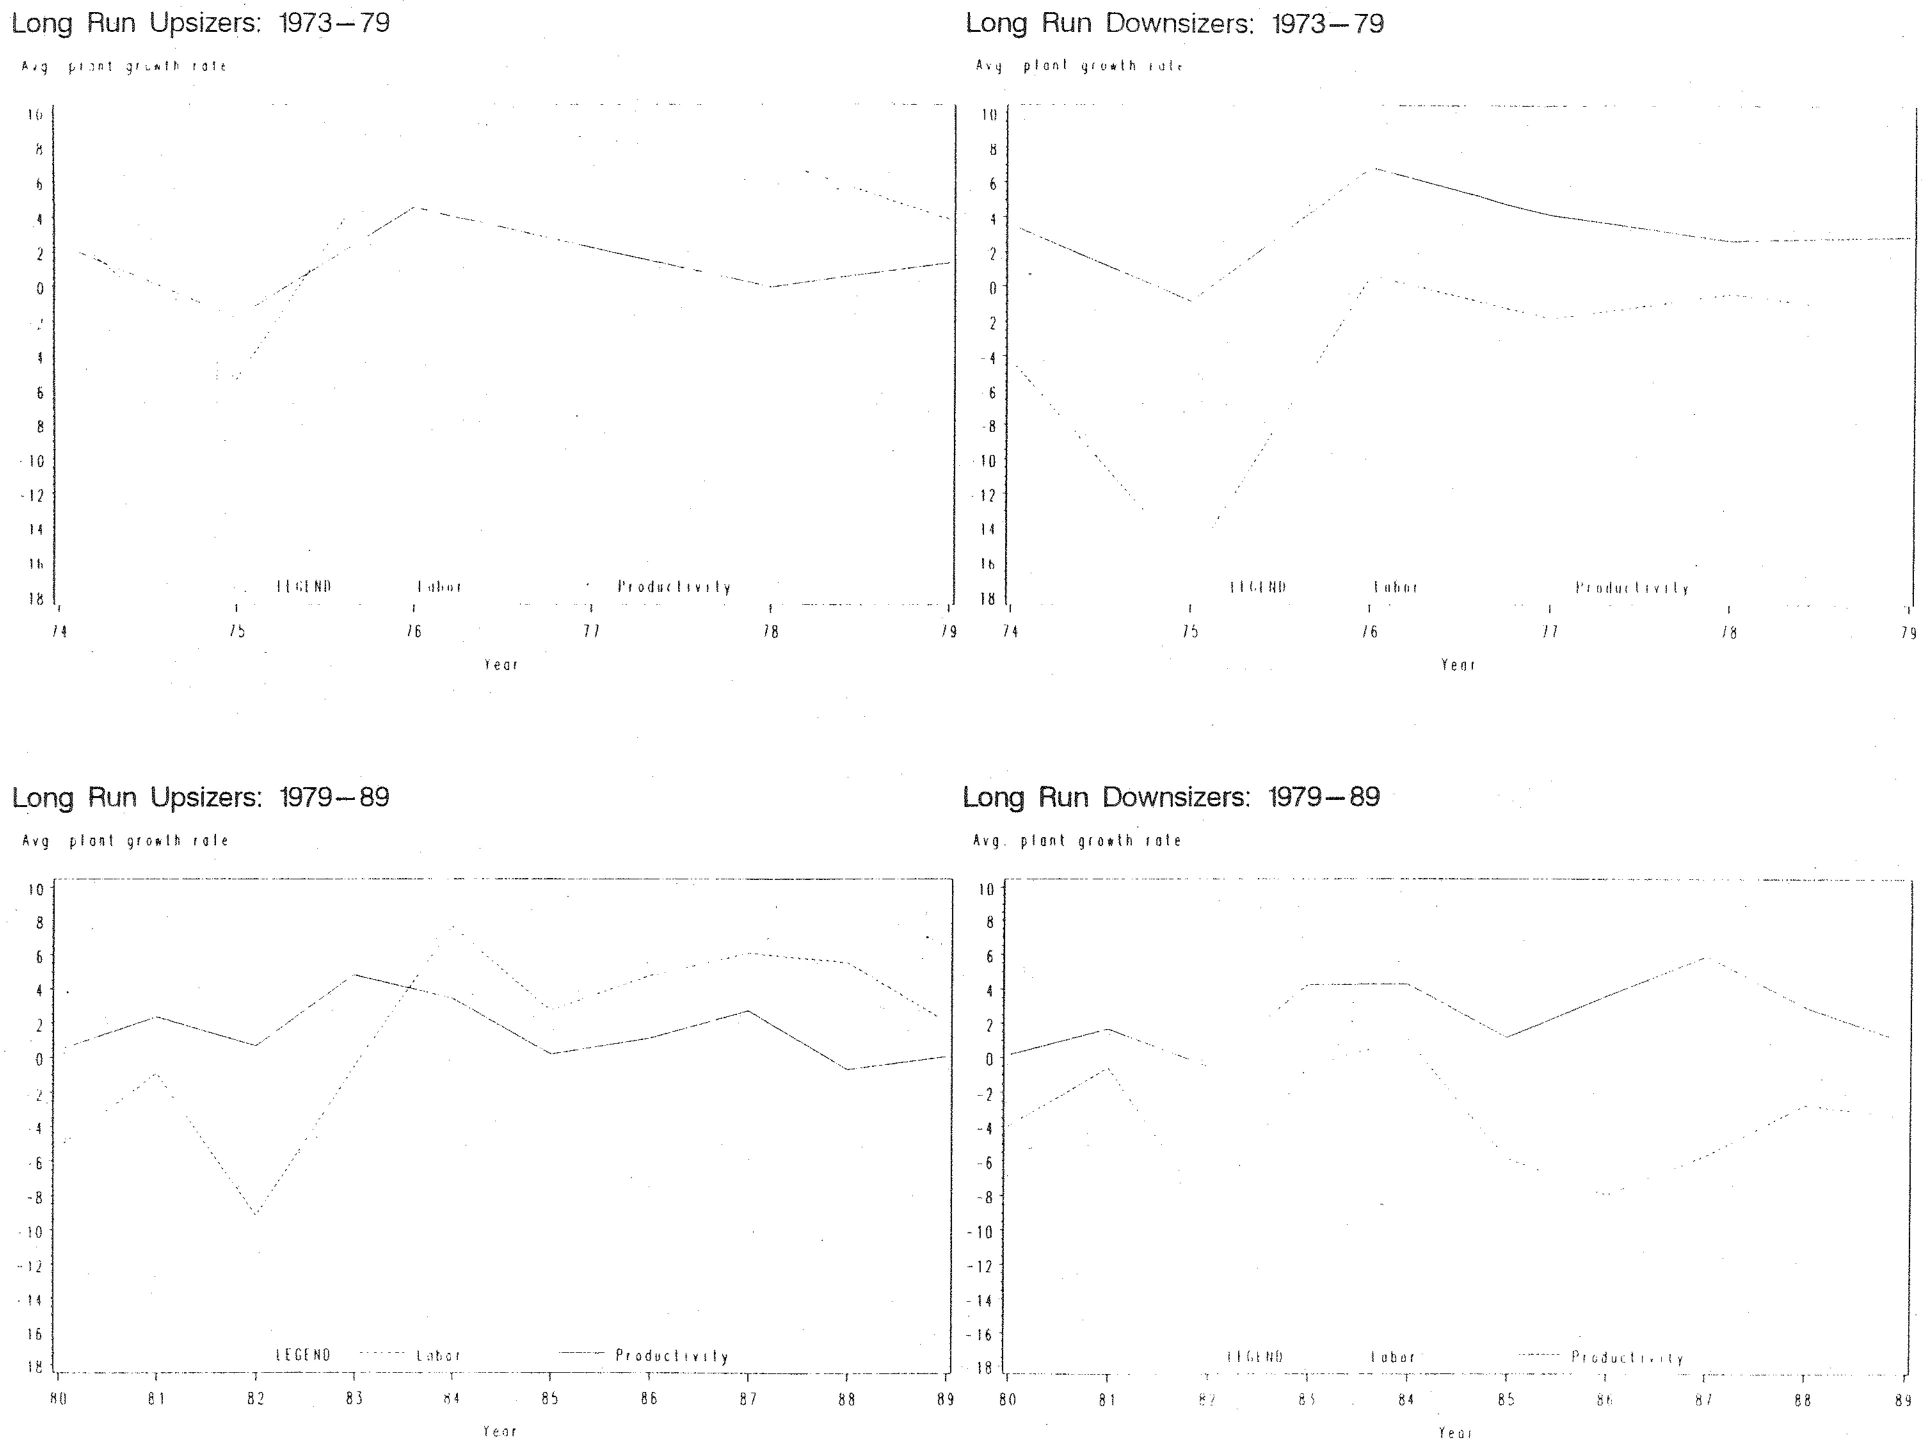

Source: Tabulations from the LRD.

Table 1: Summary Statistics on Long Run Upsizers and Downsizers

Statistic:	1973-79		1979-89	
	Long Run Downsizers	Long Run Upsizers	Long Run Downsizers	Long Run Upsizers
Mean Productivity Growth	3.2	1.4	2.4	1.5
Standard Deviation of Productivity Growth	2.5	2.2	2.1	1.7
Mean Employment Growth	-3.9	4.6	-3.9	1.3
Standard Deviation of Employment Growth	6.2	5.5	3.3	5.3

Source: Tabulations from the LRD.

Table 2: Plant-Level Cyclical Productivity

Sample Period	Dependent Variable	Coefficient Estimate on Cyclical Demand Indicator
1973-79	Output Per Hour	.235 (.015)
1979-89	Output Per Hour	.238 (.013)
1973-79	Output Per Employee	.468 (.015)
1979-89	Output Per Employee	.436 (.013)

Notes: Reported results are from bivariate regression of plant level (log) productivity on downstream cyclical indicator. Standard errors reported in parentheses.

Table 3a: Cyclical Productivity, 1979-89 (Ex-Post Definitions of Long Run Groups)

Equation	Dependent Variable	Interaction Group: Long-Run Upsizers (γ_1)		Interaction Group: Long-Run Downsizers (γ_2)	
2	Output per Hour	.066 (.021)		.346 (.017)	
2	Output per Employee	.201 (.021)		.582 (.017)	
3	Output per Hour	.390 (.422)		5.34 (.233)	
3	Output per Employee	1.81 (.420)		7.59 (.232)	
		Expansion (γ_{1E})	Contraction (γ_{1R})	Expansion (γ_{2E})	Contraction (γ_{2R})
4	Output per Hour	.022 (,030)	.109 (.030)	.321 (.024)	.370 (.024)
4	Output per Employee	.133 (.030)	.271 (.030)	.522 (.024)	.638 (.024)
5	Output per Hour	.163 (.595)	.623 (.603)	5.32 (.340)	5.36 (.324)
5	Output per Employee	1.32 (.592)	2.31 (.600)	7.27 (.338)	7.89 (.322)

Standard errors reported in parentheses.

Table 3b: Cyclical Productivity, 1979-89 (Ex-Ante Definitions of Long Run Groups).

		Interaction Group	
Equation	Dependent Variable	Long-Run Upsizers (γ_1)	Long-Run Downsizers (γ_2)
2	Output per Hour	.076 (.025)	.303 (.016)
2	Output per Employee	.206 (.024)	.528 (.015)
3	Output per Hour	4.99 (.968)	8.73 (.384)
3	Output per Employee	8.43 (.963)	13.33 (.381)

		Expansion (γ_{1E})	Contraction (γ_{1R})	Expansion (γ_{2E})	Contraction (γ_{2R})
4	Output per Hour	.030 (.035)	.125 (.035)	.278 (.023)	.327 (.022)
4	Output per Employee	.130 (.034)	.285 (.035)	.472 (.023)	.582 (.022)
5	Output per Hour	3.66 (1.38)	6.30 (1.37)	8.45 (.576)	8.96 (.526)
5	Output per Employee	6.50 (1.37)	10.33 (1.36)	12.76 (.572)	13.80 (.523)

Standard errors reported in parentheses.

Table 3c: Cyclical Productivity, 1973-79 (Ex-Post Definitions of Long Run Groups)

Equation	Dependent Variable	Interaction Group: Long-Run Upsizers (γ_1)		Interaction Group: Long-Run Downsizers (γ_2)	
2	Output per Hour	.179 (.019)		.318 (.023)	
2	Output per Employee	.389 (.019)		.584 (.023)	
3	Output per Hour	.577 (.230)		5.45 (.371)	
3	Output per Employee	2.07 (.233)		8.38 (.376)	
		Expansion (γ_{1E})	Contraction (γ_{1R})	Expansion (γ_{2E})	Contraction (γ_{2R})
4	Output per Hour	.053 (.029)	.286 (.026)	.201 (.034)	.417 (.032)
4	Output per Employee	.226 (.029)	.528 (.026)	.426 (.035)	.721 (.032)
5	Output per Hour	-.139 (.342)	1.20 (.318)	5.22 (.545)	5.66 (.514)
5	Output per Employee	.977 (.346)	3.03 (.323)	7.63 (.552)	9.06 (.521)

Standard errors reported in parentheses.

Table 4: F-tests on Coefficient Differences Between Upsizers and Downsizers

		Marginal Significance Levels		
Equation	Dependent Variable	$\gamma_1=\gamma_2$	$\gamma_{1E}=\gamma_{2E}$	$\gamma_{1R}=\gamma_{2R}$
		Panel A: 1979-89, Ex Post		
1	Output per Hour	0.0001	0.0001	0.0001
1	Output per Employee	0.0001	0.0001	0.0001
2	Output per Hour	0.0001	0.0001	0.0001
2	Output per Employee	0.0001	0.0001	0.0001
		Panel B: 1979-89, Ex Ante		
1	Output per Hour	0.0001	0.0001	0.0001
1	Output per Employee	0.0001	0.0001	0.0001
2	Output per Hour	0.0003	0.0013	0.0679
2	Output per Employee	0.0001	0.0001	0.0001
		Panel C: 1973-79, Ex Post		
1	Output per Hour	0.0001	0.0007	0.0012
1	Output per Employee	0.0001	0.0001	0.0214
2	Output per Hour	0.0001	0.0001	0.0001
2	Output per Employee	0.0001	0.0001	0.0001